

AD-A189 767

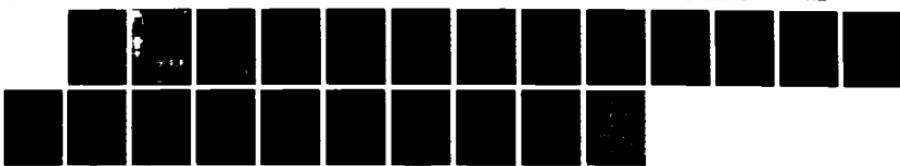
PARALLEL MEMORY ADDRESSING USING COINCIDENT OPTICAL
PULSES (U) PITTSBURGH UNIV PA DEPT OF COMPUTER SCIENCE
D CHIRULLI ET AL. DEC 87

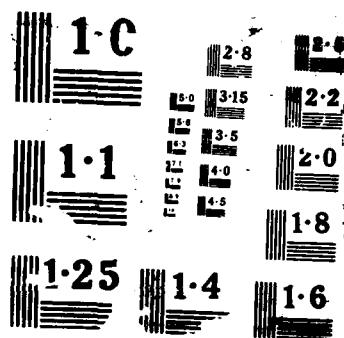
1/1

UNCLASSIFIED

F/G 12/6

NL





DEPARTMENT
OF
COMPUTER SCIENCE

AD-A189 767

Parallel Memory Addressing Using
Coincident Optical Pulses¹

Donald Chirulli
Rami Melhem
and
Steve Levitan

DTIC
SELECTED
DEC 11 1987
S D



University of Pittsburgh
Pittsburgh, Pennsylvania 15260

This document has been approved
for public release and does not
contain recommendations or conclusions
of the Defense Intelligence Agency.

87 106 138

Parallel Memory Addressing Using
Coincident Optical Pulses^{*)}

Donald Chirulli
Rami Melhem
and
Steve Levitan

To appear in a special issue of the
IEEE Computer Magazine
on
Integrated Optics for Computer and Signal Processing
(Scheduled for December 1987)

DTIC
ELECTED
DEC 11 1987
S E D
RE

⁶⁹⁵³
* This research is in part supported under ONR Contract N00014 80 C 0011.

This document has been approved
for public release and sale.
Distribution is unlimited.

Parallel Memory Addressing Using Coincident Optical Pulses

Donald M. Chiarulli

Rami Melhem

Department of Computer Science

Steven Levitan

Department of Electrical Engineering

University of Pittsburgh
Pittsburgh, PA. 15260

Accession For	
NTIS	GRA&I
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification <i>yes</i>	
By _____	
Distribution/ _____	
Availability Codes _____	
Dist	Avail and/or Special
A-1	

ABSTRACT

Common bus, shared memory, multiprocessors are the most widely used parallel processing architectures. Unfortunately, such systems suffer from a memory/bus bandwidth limitation problem. For the designer of a hybrid optical/electronic supercomputer, an immediate temptation is to replace the shared electronic bus with an optical analog of higher bandwidth. To make that replacement is only a partial solution. The true bottleneck in such systems is in the address decoding circuits of shared memory units.



In this paper we propose a new memory structure which provides for parallel access in a multiprocessor environment. The proposed system has two advantages. First, it distributes the address decoding circuitry to each of the requesting units on a common bus, thus eliminating the bottleneck of centralized decoding of encoded memory addresses. Second, it allows for parallel fetches of memory data with a level of parallelism limited only by the ratios of optical to electronic bus bandwidths and the dimensionality of the memory array. The design requires no active optical switching devices and can be implemented using the mature technologies of optical sources, waveguides, and detectors. The two-dimensional version of the device is well suited to an integrated optics implementation.

Parallel Memory Addressing Using Coincident Optical Pulses

Donald M. Chiarulli

Rami Melhem

Department of Computer Science

Steven Levitan

Department of Electrical Engineering

University of Pittsburgh
Pittsburgh, PA. 15260

1. Introduction

In a conventional electronic memory circuit, like the one shown in Figure 1, an incoming memory address is divided into row and column addresses, each decoded separately. The selected memory location is the intersection of the select lines generated by the row and column decoders. In common bus multiprocessors these decoders have traditionally been a performance limiting bottleneck. Each decoder can process only a single encoded address, thus limiting memory access to a single location. Memory interleaving techniques[1], which subdivide the memory space into regions, each in a separate memory unit, are commonly applied in an attempt to make parallel some subset of memory accesses. More recently, sophisticated cache memory systems[2] have been developed which physically reproduce portions of shared memories in a local store. Both systems have obvious limitations. Interleaved systems impose an ordering in which parallel accesses to a shared memory may be made, and cache memories rely on the locality of memory references for each processor and require a large overhead to support cache coherence.

Our solution is to distribute the address decoding function to the requesting devices, thereby breaking contention for monolithic address decoders. This solution requires the abandonment of conventional encoded addresses as a mechanism for conserving bus bandwidth. Rather, we will use the high bandwidth of optics to time multiplex fully decoded addresses into an optical "select" pulse train. Using a technique based on the coincidence of optical pulses, the optical select pulse train can be directly applied to a memory array to address one or more cells. Effective parallelism is possible in this system because of the differential between optical and

This research has been supported, in part, under Office of Naval Research contract N00014-85-K-0339.

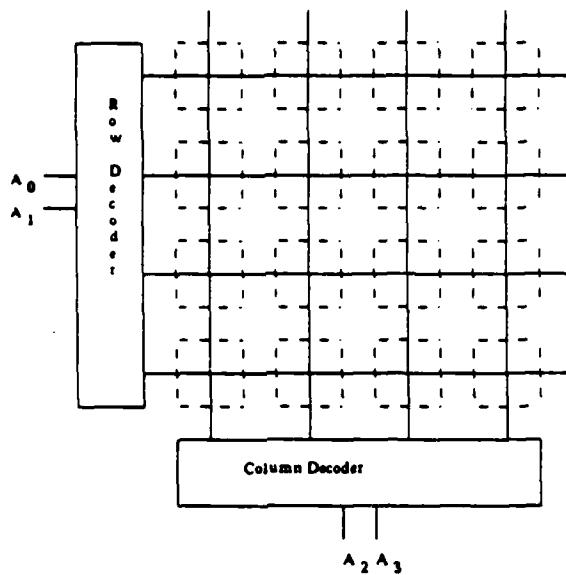


Figure 1: Conventional Electronic Memory Structure

electronic bandwidth. Within a single electronic memory access cycle, N parallel memory references are possible, where N is limited only by the ratio of optical to electronic bandwidths. For a fixed bandwidth ratio, the size of memory which can be constructed is further determined by the dimensionality of the memory structure. The technique requires no active optical or electro-optical switching devices. It uses only the mature technologies of optical sources, waveguides, and photodetectors. In the two-dimensional form, the system is well adapted to an integrated optics implementation[3, 4].

The addressing mechanism, which we call "optical pulse delay modulation", is based on the use of time delays between optical pulses. The optical pulses are propagated through waveguides in several directions through the memory array. By appropriately adjusting the delays, these pulses can be made to coincide at specific memory cells. This coincidence is detected by photodetectors at the addressed locations, thereby selecting those locations for memory access.

Our primary interest is in the application of this addressing mechanism to two-dimensional, multi-ported, memory modules. Such structures are composed of horizontal and vertical waveguides with n memory cells located at the intersecting points. With proper cell layout, up to \sqrt{n} memory cells may be accessed concurrently by sending a sequence of pulses in the horizontal and vertical waveguides. In a multiprocessor environment, a sequence of pulses, each corresponding to a distinct memory reference, is generated by independent address decoders located at each of the processing units. Therefore, the address decoding function is completely distributed to the requesting processors and there is no

address decoding circuitry at the memory unit.

This paper is organized as follows, in section two we introduce coincident pulse addressing by an example using a one dimensional memory structure. Included in this section is a discussion of techniques for selecting memory locations, generating parallel addresses, and memory data output. In section three we expand the technique to a more realistic two-dimensional structure and address such issues as memory access conflicts, complexity, and organization. Suggested systems for memory write control are presented in section four. In the concluding section, we discuss higher dimensional structures, implementation issues, and the research issues which must be addressed to physically realize such a system.

2. A One Dimensional Memory Array

In this section, we introduce the technique of pulse delay addressing using an one-dimensional memory array as an example. This example is not ideal because both the hardware complexity and access time grow in proportion to the size of the memory. This is not the case for the higher dimensional structures presented in later sections.

2.1. Optical pulse delay addressing

As shown in Figure 2, a memory module is composed of n cells C_1, \dots, C_n , each storing one bit of information. The select signal for each cell C_k , is an electronic pulse at the output of a photodetector D_k . The photodetector generates the logical OR of two incident optical signals, denoted in Figure 2 by s_1 and s_2 .

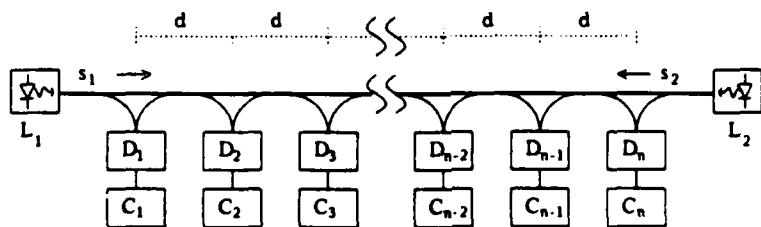


Figure 2: Linear Memory Structure

The signals s_1 and s_2 travel in opposite directions along an optical path, which can be either an optical fiber or a planar waveguide in an integrator optical device.[5] Photodetectors are placed at fixed distance intervals d along the optical path, and two laser diodes, L_1 and L_2 , are coupled to each end. Both laser diodes are normally on and the circuits of all detectors normally generate a logic one.

Assume that two dark pulses of duration τ are transmitted, one from L_1 and the other from L_2 at times t_1 and t_2 , respectively. These pulses represent "dark" spots

propagating at speed c_g (the speed of light in the waveguide). By carefully selecting the delay between t_1 and t_2 , the two dark spots can be made to meet at exactly one detector. This detector will then turn off, generating a logic *zero* of duration τ . The distance d between any two detectors is chosen to be equal to $d = \tau c_g$, the propagation distance corresponding to the pulse duty cycle. The delay $t_1 - t_2$ is also chosen such that it is an even multiple of d . More specifically, if

$$t_1 - t_2 = (n - 1 - 2(k - 1)) \tau \quad (1)$$

then, the two dark spots will meet at detector D_k , thus addressing cell k . For example, when $n = 5$, if L_2 generates its dark pulse 2τ seconds before L_1 generates its pulse, then (1) gives $k = 2$, that is the two pulses meet at D_2 . Similarly, if L_2 generates its pulse 2τ seconds after L_1 generates its pulse, then the two pulses meet at D_4 . Clearly, the middle cell is chosen by generating the two pulses simultaneously, that is having $t_1 = t_2$. Therefore, the address of the cell is encoded using the delay $t_1 - t_2$. In this view, the pulse generated by L_1 is treated as the reference pulse and the pulse generated by L_2 becomes a select pulse. In the remaining discussion, the names t_{ref} , L_{ref} , t_{sel} , and L_{sel} , will refer to t_1 , L_1 , t_2 , and L_2 respectively.

The memory access time is determined by the maximum delay needed to address any cell in the array. From (1), it is clear that for $k = 1, \dots, n$, we have

$$-(n - 1) \tau \leq t_{ref} - t_{sel} \leq (n - 1) \tau \quad (2a)$$

from which we find that the memory access time, σ , is given by

$$\sigma = 2n\tau \quad (2b)$$

Note that equation (2a) indicates that the select pulse occurs within $n\tau$ before or after the reference pulse.

The parallelism in this addressing scheme comes from the fact that within time σ , it is possible to address more than one cell by sending a series of pulses from L_{sel} , one for each memory reference. Each of these pulses will intersect with the reference pulse at the desired detector. In other words, parallel memory references are positionally distinguishable in a pulse train generated by a series of select pulses.

In the next two sections we will describe how this memory can be incorporated into a shared memory multiprocessor. Two design issues at the interface between the electronic processing units and the optical system must be resolved. First, a system for generating a series of optical pulses, corresponding to the select pulse train must be specified. Second, the memory must allow data stored in the referenced locations to be returned to the requesting processors within a single

processor-memory cycle.

2.2. Address Generation

One of the major advantages of the proposed memory organization over conventional systems is the removal of the address decoding function from the memory unit, and the distribution of this function to the requesting processors. More specifically, each of the processors is assumed to generate normal encoded addresses when referencing the shared memory. These addresses are decoded locally by an address decoder attached to each processor. The decoded addresses are electronically OR'ed onto a select bus common to all of the processors. In the one-dimensional case the select bus consists of n lines, each controlling a laser diode pulser (see Figure 3). As will be explained in Section 3, the size of the select bus in the two-dimensional case reduces to a more manageable $2\sqrt{n}$ lines controlling $2\sqrt{n}$ laser diode pulsers.

Returning to the linear case, the optical pulse train containing the memory requests is generated by $2n$ laser diode pulsers which are spaced at incremental distances d in optical path length from the memory. In order to reconcile the optical to electronic bandwidth difference, a single edge in the electronic time-base (denoted *sync 1* in Figure 3) controls the activation of all the pulsers such that all the optical pulses are generated simultaneously. If the duration of each optical pulse is equal to τ , then the select pulse train will be confined to $2n$ time slots, each with duration τ . Since proper time multiplexing of select pulses is not possible with the select pulsers constantly on, the n select pulsers connected to the electronic select bus are separated by n pulsers modulated by a fixed *one* signal (see Fig. 3). Thus all the slots will contain an optical pulse except those slots corresponding to requested memory addresses. More specifically, a dark spot (no pulse) at slot $2i - 1$ corresponds to a request for memory location i .

In addition to the select pulse train, a dark reference pulse must be generated. Alternatively, the reference pulse is a series of $2n$ light pulses, with a single dark pulse at position n . Such a pulse train may be produced by a single laser diode which is normally *on* and which is *pulsed off*, for duration τ , upon the reception of the synchronization pulse *sync1*. For the dark pulse in the reference train to coincide with pulse slot n in the select pulse train, the optical path from the reference diode to the memory should be equal to the optical path from the diode generating the middle pulse in the select train to the memory.

The above address generation scheme is crucially dependent on the simultaneous pulsing of several laser diodes. In an integrated optics environment, such synchronization problems can be avoided by using a single optical pulse source of

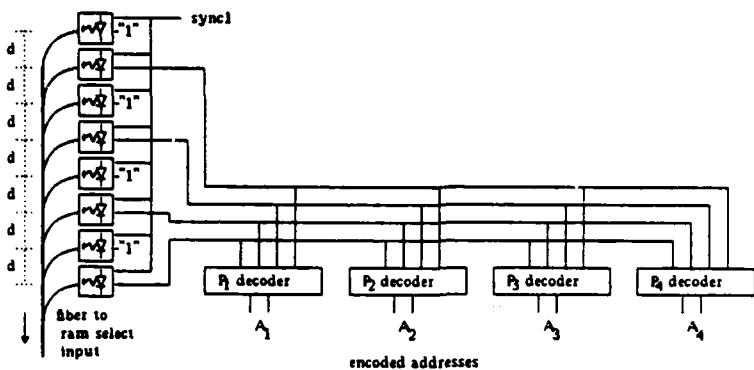


Figure 3: Distributed Address Generator

duration τ , and a series of electro-optic switches. In such an implementation each laser diode in Figure 3 is replaced by a directional coupler which "couples in" the pulse at various optical path lengths. Synchronization problems are replaced in this system by a new set of issues relating to optical power distribution. We will discuss this and other issues relating to an integrated optics implementation in later sections.

2.3. Parallel Memory Read

Another issue at the interface between the electronic processing units and the optical system is a mechanism for returning the electronically stored data from the memory to the processors. The data is returned, on an optical bus, in a pulse train which consists of n slots, one for each memory location. Therefore, parallel accesses are positionally distinguishable in the read pulse train. When this pulse train arrives at a processor which has issued a read request for the i^{th} position of the store, this processor will find the requested data in the i^{th} slot of the data pulse train.

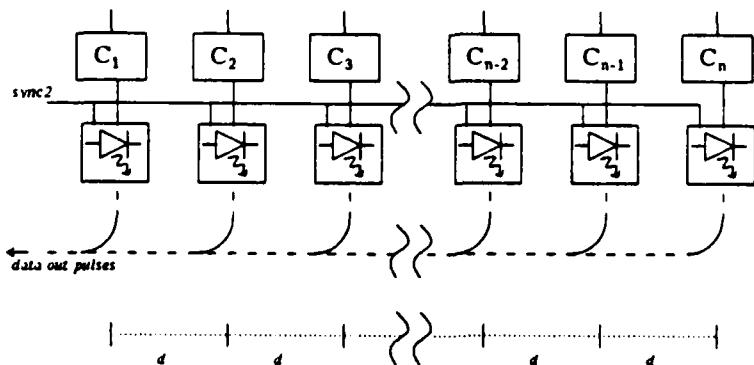


Figure 4 - The generation of the data pulse train

One method for generating the read data pulse train is to use a structure similar to the one depicted in Figure 4. In this structure, n laser diodes are placed on the optical data bus separated by an optical distance d . When a specific memory location k is addressed, an electronic signal is generated from the photodetector as described in Section 2.1. The data at location k is assumed to be stored electronically and is used to modulate the k^{th} laser diode only if location k is addressed. A synchronization signal, sync 2, is used to synchronize the output of light (positive) pulses of duration τ from the selected memory locations which store a *one*. The difference in the optical path lengths between the laser diodes ensures the correct generation of the data pulse train. A similar technique is used to demultiplex the pulse trains by using detectors at fixed distances d , and latching the pulse train at each processor interface.

3. A Two Dimensional Memory Structure

Using the above mechanism, it is possible to address all of the n memory locations in one processor-memory cycle. For the one-dimensional case, this represents a sequential read of the entire store, and requires a ratio of electronic to optical time bases equal to the size of the memory. Even for the most optimistic assumptions about achievable optical pulse widths, this structure will be inadequate and wasteful. However, by applying a similar addressing mechanism to two-dimensional memory arrays, the required ratio of electronic to optical time bases reduces to \sqrt{n} , where at most \sqrt{n} memory locations may be addressed in one cycle. This allows for the construction of reasonable size shared memories.

3.1. Coincident Wavefront Addressing

In the two-dimensional case we generalize the propagation of "dark spots" in one dimension to the propagation of linear dark wavefronts moving through a series of parallel waveguides. Hence, the method of addressing a location by programming the intersection of two dark spots may be generalized to addressing a location in a two-dimensional array by programming the intersection of three dark wavefronts. The literature on systolic and wavefront arrays (see for e.g.[6, 7]) suggests many possible ways for propagating and programming the intersection of wavefronts. In this section, we present a simple propagation scheme which may be used in 2-dimension memory addressing.

Consider 2-D memory arrays similar to the one shown in Figure 5. An array of size n is composed of $\sqrt{n} \times \sqrt{n}$ photodetector/cell units $DC_{i,j}$, $i, j = 1, \dots, \sqrt{n}$, separated by a distance $d = \tau c_g$ in both the vertical and the horizontal directions. The structure of a DC unit is identical to the linear example, except that, the input

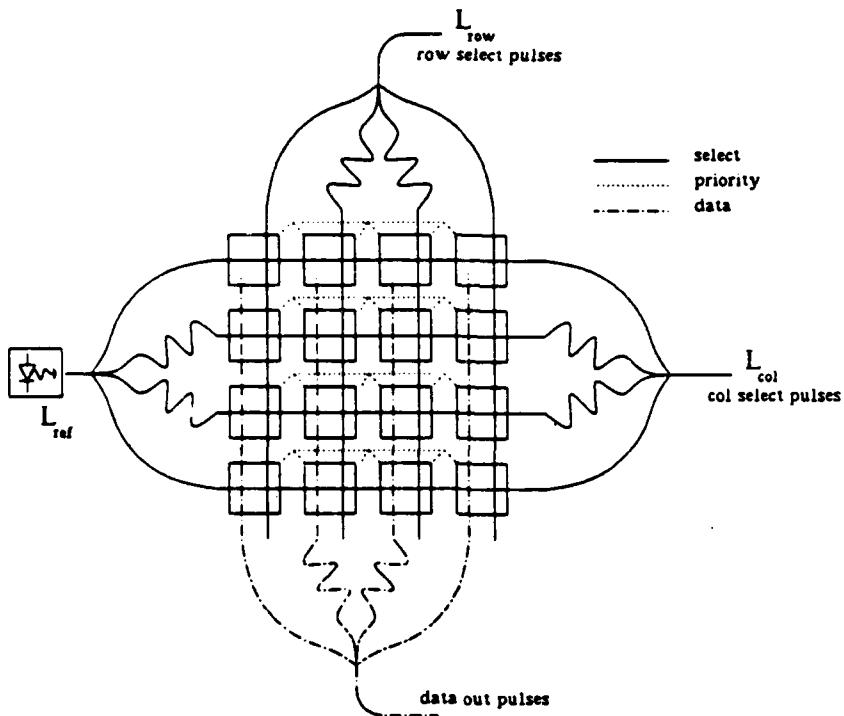


Figure 5: Two Dimensional Memory Structure

to the photodetector generates the logical OR of three optical signals. Specifically, a dark reference wavefront generated by the reference diode L_{ref} , a select pulse train generated from a distributed set of column address decoders, L_{col} , both traveling horizontally in opposite directions, and a select pulse train generated by a distributed set of row select decoders, L_{row} , traveling vertically.

The optical signal generated by each source is decoupled from the source fiber by a "squid" type connection into \sqrt{n} signals that travel through the array in parallel waveguides. Since the optical path length of all legs in the squid will be equal, the wavefront will arrive at all locations in a single row (or column) simultaneously. For example, an optical pulse generated by L_{ref} and directed horizontally through the array will be simultaneously incident at locations $DC_{j,i}$, $j = 1, \dots, \sqrt{n}$, in column i . Similarly, any pulse generated by L_{col} , will arrive at all the cells in a specific column simultaneously, and any pulse generated by L_{row} will arrive at all the cells in a specific row, simultaneously.

In order to derive the equations that govern the intersections of three wavefronts, assume, as in the case of the linear array, that all three laser diodes are on and that L_{ref} generates a dark pulse of duration τ at time t_{ref} . Also, L_{col} and L_{row} generate dark pulses at times t_{col} and t_{row} , respectively. If the timing of L_{col} is such that

$$t_{ref} - t_{col} = (\sqrt{n} - 1 - 2(j-1)) \tau \quad (3)$$

then, the two dark wavefronts generated by L_{ref} and L_{col} will meet at column j of the array. In order to select a particular cell $DC_{i,j}$ in that column, the third dark wavefront, namely the one generated by L_{row} , should be crossing row i when the other two wavefronts meet at column j . This may be accomplished by timing L_{row} such that

$$t_{ref} - t_{row} = (j - i) \tau \quad (4)$$

In other words, to address a certain memory location i, j , the column number j is encoded as $t_{ref} - t_{col}$ and the difference, $j - i$, between the column number and the row number is encoded as $t_{ref} - t_{row}$. From (3) and (4), it may be shown that

$$-(\sqrt{n} - 1) \tau \leq t_{ref} - t_{col} \leq (\sqrt{n} - 1) \tau$$

and

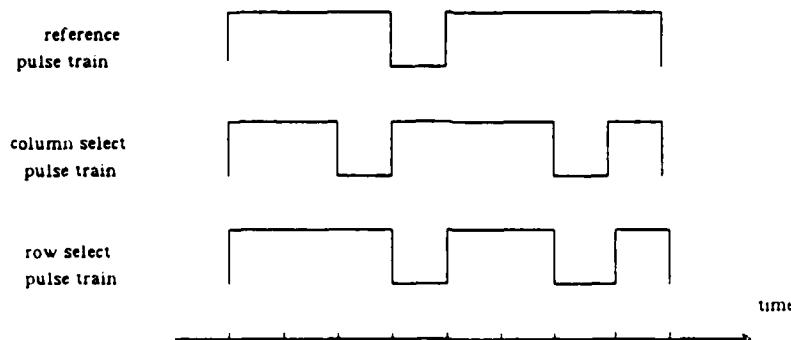
$$-(\sqrt{n} - 1) \tau \leq t_{ref} - t_{row} \leq (\sqrt{n} - 1) \tau$$

and hence, the memory access time, σ , is

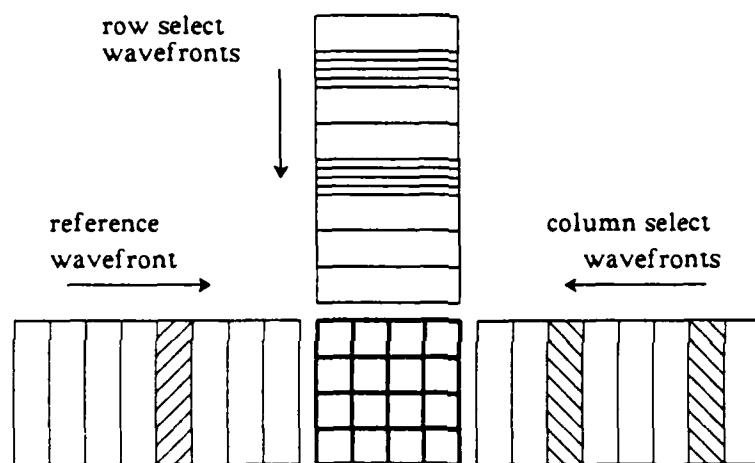
$$\sigma = 2 \sqrt{n} \tau \quad (5)$$

As in the linear case, parallel accesses are possible by generating multiple pulses in the row and column select signals. For example, Figure 6(a) shows the pulse trains for the selection of memory locations (2,2), (1,4) and (4,4) in the 16 location memory array of figure 6(b). For these three locations, $t_{ref} - t_{col}$ should be equal to -1, 3 and 3, respectively, and $t_{ref} - t_{row}$ should be equal to 0, 3 and 0, respectively. The locations of the wavefront resulting from these trains at times 0, 5τ and 7τ are shown in Fig 6(b/c/d). From these figures, it is clear from the intersection of the dark fronts that location (2,2) is selected at time 5τ , and locations (1,4) and (4,4) are selected at time 7τ .

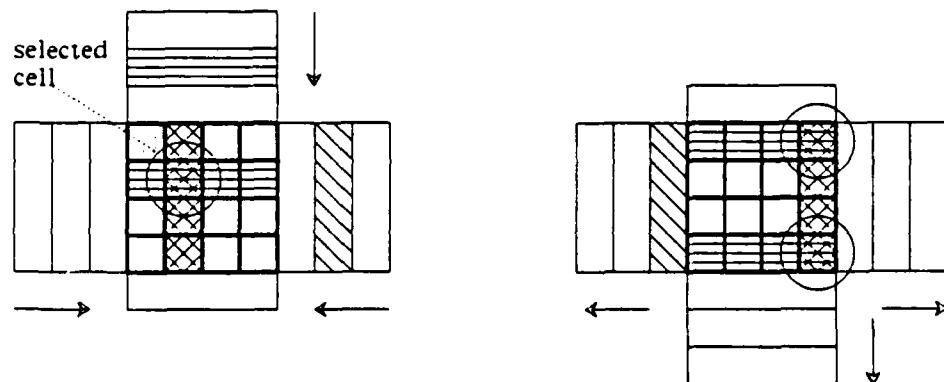
Using the above scheme, it is possible to encode the addresses of all of the n memory cells in the column and row pulse trains during a single memory access cycle. However, for a time multiplexed memory read such as the one proposed in section 2.3, the length of the return data pulse train, and hence the total read time, would grow linearly with memory size. To prevent this, and to facilitate a pipelined implementation, the maximum length of the read pulse train is chosen to be $2\sqrt{n}$, the length of the select pulse trains. This is actually an advantage of the two-dimensional structure. More specifically, in the one-dimensional case, the memory cycle time was directly proportional to the size of the store. Each cycle needed to provide an optical time-base slot for each location. In the two-dimensional case, access time, and the possible number of parallel accesses, is proportional to the square root of the number of locations in the store. This is a far



(a) select and reference pulse trains



(b) wavefronts at time 0



(c)(d) wavefronts at times 5τ and 7τ

Figure 6 - Programming the intersection of wavefronts

more realistic scenario for constructing a shared memory multiprocessor system. The price paid for this reduction in access time is the potential for conflicts in parallel memory accesses. In the next section, we suggest a technique which allows for at most \sqrt{n} distinct memory references per cycle.

3.2. Resolution of conflicts in memory access

There are two policies that may be applied to limit the number of memory references during a given cycle to \sqrt{n} . The first policy is to allow only \sqrt{n} addresses to reach the memory module during the cycle, and the second is to allow as many as n addresses to reach the memory but return only the contents of \sqrt{n} of these addresses. The first policy requires active optical switching devices to resolve conflicts in the incoming select pulse trains. To avoid the need for such devices, we choose to implement the second policy, which allows full addressing and prioritizes referenced locations for conflict resolution.

The same data collection circuitry used in the linear array is used to collect data for each column of the two-dimensional array. One waveguide is dedicated to the collection of the contents of the addressed cells in each column. The signals in the \sqrt{n} waveguides are merged into a single wave guide (denoted "data out pulses" in Figure 5) which returns the data to the processors. The lengths of the waveguides are adjusted such that the optical paths between any two cells i, j and i, j' in the same row, i , and the merging point are equal.

With this collection mechanism, the content of any referenced cell i, j in column j will appear in the $(\sqrt{n} - i + 1)^{th}$ time slot on the waveguide of column j . However, when the \sqrt{n} pulse trains corresponding to the \sqrt{n} columns are merged in the data out waveguide, the data produced by any two cells i, j and i, j' in the same row and different columns will collide. Since we have elected not to provide a mechanism for preventing conflicting addresses in the input pulse train, conflict resolution must be built into the memory array. That is, requests for memory references in the same row may be allowed, but only one request should be satisfied. Two problems arise for such a system. First, a mechanism to allow only one cell per row to output its data must be devised. Second, a method is needed for announcing which of the conflicting requests have been satisfied.

We start by resolving the second issue. The discussion to this point has assumed a single bit memory. In fact, memory locations will contain an entire word of W bits, stored electronically and returned in parallel on W optical data out lines. If each memory location is tagged with its column number, then each processor may read the column address along with the data and use the data only if the address coincides with its request. If not, the processor must re-issue the memory request.

For a memory of size n , $\log(\sqrt{n})$ tag bits are needed, which increases the number of bits stored in each memory location from W to $W + \log(\sqrt{n})$.

The first of the conflict resolution issues is the more difficult. To prevent bus contention between conflicting requests, we have chosen a priority system based on an optical priority chain, like the one shown in Figure 7. This figure shows a single row of a two-dimensional array. The optical distance from each cell in this row to the data output waveguide is equal and hence any parallel accesses within this row will conflict. In order to avoid this conflict, only one of the optical sources along this row may be allowed to generate data. In Figure 7a, the horizontal waveguide connecting all cells in the row forms an optical priority chain to resolve these conflicts.

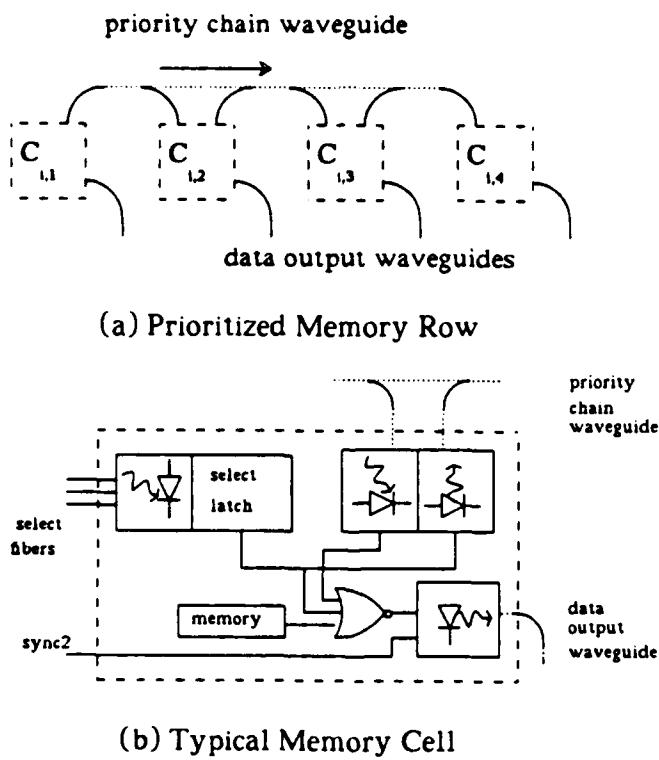


Figure 7 - Conflict Resolution Strategy

Figure 7b shows a diagram of a typical memory cell. The optical OR output from the pulse sensing photodetector sets the select latch. This output gates the contents of the memory cell through the three input electronic output-control-NOR gate. The third input to this NOR gate is the priority control. This signal is the output of a photodetector which senses select signals for higher priority cells indicated optically on the priority chain waveguide. The local select signal is also used to turn on a laser diode, which injects light into the priority chain waveguide to indicate its selection to lower priority cells. Finally, the synchronization signal, sync 2,

ensures that all data out pulsers are activated simultaneously.

From the above description, it is clear that every memory read cycle is divided into three stages. In the first stage, the select pulse trains are generated and propagated through the memory array. The minimum time required to complete this stage is equal to $2\sqrt{n} \tau$.

In the second stage, the read operation is propagated through the memory cell electronics. Note that in Figure 5, the priority chain waveguides are parallel to the reference pulse waveguides. In this arrangement, the wavefront which encodes current priority propagates through the priority chain waveguides during the first stage of the pipeline. The priority is delayed relative to the reference pulse by time

$$t_{pd} = t_s + t_l,$$

where t_s is the switching delay of the detector/latch circuit and t_l is the turn on time for the priority-out laser diode. At each cell in the rightmost column of Figure 5, which are the last cells to see the reference pulse and thus the lowest in the priority chain, the priority input signal arrives at the output control gate at time $t_{pd} + t_d$ relative to the arrival of the reference pulse. (t_d is the response time of the priority-in detector.) Since the select signal arrives at the output control gate in time t_s , the critical timing path for the second stage of the pipeline is

$$t_{pd} + t_d + t_g,$$

where t_g is the output control gate switching time. By noting that t_l and t_d must be less than or equal to the pulse width τ , we can place the lower bound on second stage pipeline delay at

$$t_s + t_g + \tau$$

Finally, in the third stage of the pipeline, data is returned from the memory to the requesting processors in a pulse train of \sqrt{n} bits. Thus, the minimum time required for the third stage is $\sqrt{n} \tau$. Assuming that the memory size and the ratio of electronic to optical bandwidths is sufficient to satisfy

$$t_s + t_g < (2\sqrt{n} - 1)\tau,$$

the longest of the stages is the first. Using such a three stage pipeline the total memory cycle length is then $6\sqrt{n} \tau$. Since we are accessing the memory in a pipelined fashion, and each stage can process \sqrt{n} references, the effective memory bandwidth limit is $1/(2\tau)$ words/second.

Finally, it should be mentioned that it is possible to support $2\sqrt{n}$ memory references per cycle, rather than merely \sqrt{n} , by rearranging the data collection waveguides of Figure 5. If the data collection waveguides are run diagonally, $2\sqrt{n}$

waveguides can be accommodated at the price of more complex and unevenly distributed conflict resolution scheme.

3.3. Organization of Memory Modules

In the above discussion, we have concentrated on select and read mechanisms assuming a one bit memory. In an $n \times W$ bit memory module we would reproduce W copies of the memory cell, the output-control-NOR gate, and the data output pulser of Figure 7b. For reading, only one decoder-select plane, and one priority chain waveguide is necessary for each of the $W + \log(\sqrt{n})$ bit words in an n word array. The only optical system which must be scaled with the number of bits in the word is the read output pulsers. In the same manner as an electronic memory, we must provide a separate return path for each bit.

4. Memory Write Control

For a conventional memory, support for write operations would require an additional control signal and a secondary data path for incoming data. Merely providing these additional signals in a parallel memory will not be adequate, since the additional issue of resolving mixed parallel reads and writes must also be resolved. There are three possibilities for resolving this issue which trade off write access time for optical circuit complexity.

- 1) *Exclusive write*: in this solution we eliminate the possibility of mixed read/write operations and conflicting write operations. By implementing an external arbitration mechanism, we allow only exclusive write access to the memory. Once a single processor is selected, it can perform writes to the memory using conventional electronics. This system requires no additional optics at the cost of non-parallel writes. If the ratio of writes to reads for the shared memory structure is relatively low, then exclusive write access may represent a viable low complexity alternative.
- 2) *Full parallel read/write*: for full parallel optical writes, it is not actually necessary to provide a separate optical write data bus. Instead, by combining control and data information, in lieu of a data bus, it is possible to provide fully parallel non-conflicting writes to any of the n locations in the store.

In this technique, each memory bit sees two bits of select information in each cycle, one bit from the read select optical circuit already described, and a second from a *per-bit* local copy of that selection circuit used for write control. These two bits encode four states: read, write a zero, write a one, and do nothing (no select). Thus, by reproducing at each bit in the word a second address selection structure, and by judicious selection of code assignments for the four

states, the cost of this system is limited to the addition of one optical selection plane for each bit.

This technique allows fully mixed read and parallel writes. Any processor can write to any word with no conflict restrictions on the rows or columns of write addresses. However, there is no conflict resolution mechanism to prevent two processors from writing to the same address. As with a multi-ported electronic memory, such operations would generate unpredictable results. It is assumed that these mutual exclusion issues would be addressed by appropriate software.

- 3) *Bit Serial Parallel Write:* At a cost of lengthening the overall write time, it is possible to reduce the overhead for the full parallel read/write system to one additional optical selection plane by using a bit serial approach. In this system, a designated cycle initiates an W bit serial write. The select signals generated by the read selection circuit during this cycle are latched separately, and held for the duration of the serial write. Each subsequent cycle uses the write selection circuit to serially process each bit in the word. Meanwhile, parallel reads are still possible, concurrent with the serial write. A global counter/decoder circuit, to define the "current-bit" to be one of the W data bit planes of the memory, is necessary. It is the only additional overhead for this system.

5. Extensions and Future Research

We have concentrated in this presentation on the details of a two dimensional memory array because of its suitability to integrated optical implementations. There is nothing in the design that prevents the linear wavefronts in two dimensional arrays from being generalized to planar waves in three dimensions arrays. In general, for m -dimensional memory arrays, the access time is reduced to the m th root of the memory size. Hence, for a fixed bandwidth in the electronic system, the bandwidth requirements for the optical system are substantially reduced. This reduction is gained at the price of a corresponding reduction in the number of locations which can be referenced in parallel during a single electronic cycle. Like the access time, this number is also reduced to the m th root of the memory size.

Extensions of this technique can also be applied to other two-dimensional switching structures. For example, the control information for a crossbar switch can be encoded into pulse trains and used to address specific switches in the crossbar. Pulse delay encoded control information can be prepended on incoming packets to combine routing information and data without the need for optical-optical switching devices.

The question is: Can such a memory be built, and will its size and performance make it feasible for integration in computer systems of the next decade and

beyond? The following issues must be examined.

- 1) **Scalability:** The scale of the physical device is directly related to the optical pulse width. To minimize the physical spacing of detectors, very short pulses are required. For instance, to reduce detector spacing to a scale that will allow monolithic integration, pico-second pulse widths will be required. Specifically, 1ps pulses will allow a detector spacing of 200 to 300 μ meters depending on the refractive index of the optical medium. Current commercial technology for pulsing laser diodes in discrete devices provides for pulses on the order of 100ps. Recent research has produced optical pulses as short as 8 femto-seconds[8, 9]. In such research, the common technique for pulse duration measurement is to split the pulse into two optical paths and detect coincidence when the paths are recombined at varying optical path lengths. Based on this trend we expect that the necessary pulse widths for an integrated optical implementation will be available in the near future. Meanwhile, we are currently using commercial discrete devices and optical fibers to examine scalability issues.
- 2) **Detection limits:** A second limit on usable pulse duration is the issue of detector technology for coincident light pulses. A two dimensional memory requires that as many as three dark pulses are to be detected as they overlap. Even assuming the existence of photo-detectors of sufficient bandwidth for single pulse detection, what degree of overlap is required to generate the optical OR of three pulses? Extending to multi-dimensional structures, fan-in limitations of this type become even more critical.
- 3) **Fabrication vs. physical limits:** As the system scales down in size, and up in speed, what limits will be reached first, the fabrication limits of the technology or the physical limitations of the optical systems?
- 4) **Clocking issues:** While there is little doubt that sufficiently narrow pulses can be generated by the electro-optics, the precision to which multiple pulses can be synchronized is an important question. Two coincident pulses must be timed to arrive with a precision of +/-10% of their pulse width to allow at least 80% overlap. This means that the electronic components must gate the optical signals with a constant delay that is precise to the optical time-base. Clock distribution issues have been studied extensively in both electronic and optical domains. We believe that the required precision can be achieved by electronic circuitry. As an alternative, optical clock distribution techniques such as those proposed by Clymer and Goodman[10] can be applied.
- 5) **Select Latch Response:** The electronic latches at the detector sites must respond to the selection pulses from optical detectors. These pulses will be no longer than the duration of the coincident pulses. This is a limitation on the speed of

the optical system.

6) **Waveguide Decoupling Technology:** In a two-dimensional design it is necessary to split incoming optical signals into parallel row and column waveguides. Detectors at each intersection must couple out sufficient optical power for detection, without significantly degrading the optical signal. Such highly asymmetric single-mode directional couplers have been developed for optical fiber and are commercially available[11]. Several other techniques for low power output coupling have been examined by Jackson et al[12]. Further work is needed to apply these techniques in an integrated optic environment.

In summary, we have presented a system which distributes the address decoding function to the requesting units on an optical bus. In this system, addresses become optical pulse trains, and by arranging the optical paths we provide a selection mechanism based on the coincidence of these pulses. In the coming year we plan to begin construction of a 64×16 bit register file based on this research and using discrete optical devices and fiber waveguides. This register file will be used as shared memory in a prototype eight node multiprocessor.

References

1. Chen, T. C., "Overlap and Pipeline Processing," in *Introduction to Computer Architecture*, ed. Harold Stone, Science Research Associates, 1975.
2. Briggs, F. and K. Hwang, *Computer Architecture and Parallel Processing*, McGraw Hill, 1984.
3. Neff, J., "Major Initiatives for Optical Computing," *Optical Engineering*, vol. 26, no. 1, January 1987.
4. IEEE Spectrum, *Special Issue on Optical Computer*, 28, August 1986.
5. Lee, D. L., *Electromagnetic Principles of Integrated Optics*, John Wiley and Sons, New York, 1986.
6. Kung, H. T., "Why Systolic Architectures," *Computer*, vol. 15, pp. 37-46, 1982.
7. Kung S. Y., Arun, K., Gal-Ezer, R., and Rao, B., "Wavefront Array Processor: Language, Architecture, and Applications," *IEEE Transactions on Computers*, vol. 31, no. 11, pp. 1054-1066, 1982.
8. Shank, C. V., "The Role of Ultrafast Optical Pulses in High Speed Electronics," in *Picosecond Electronics and Opto-electronics*, ed. Morou, G., Bloom, D., and Lee C., Springer Verlag, 1985.

9. Fujimoto, J. G., Weiner, A. M., and Ippen, E. P., "Generation and Measurement of optical pulses as short as 16 fs.," *Applied Physics Letters*, vol. 44, no. 9, May, 1984.
10. Clymer, B. D., and J. W. Goodman, "Optical Clock Distribution to Silicon Chips," *Proceeding SPIE*, vol. 625, 1986.
11. Tekippe, V. J. and W. R. Wilson, "Single-Mode Directional Couplers," *Laser Focus*, May 1985.
12. Jackson, Kenneth P., et. al., "Optical Fiber Delay-Line Signal Processing," *IEEE Transactions on Microwave Theory and Technologies*, vol. MTT-33, no. 3, March 1985.

END

DATE

FILMED

APRIL

1988

DTIC